



Article scientifique

Article

2024

Published version

Open Access

This is the published version of the publication, made available in accordance with the publisher's policy.

Thermodynamic computing via autonomous quantum thermal machines

Lipka-Bartosik, Patryk Jan; Perarnau Llobet, Marti; Brunner, Nicolas

How to cite

LIPKA-BARTOSIK, Patryk Jan, PERARNAU LLOBET, Marti, BRUNNER, Nicolas. Thermodynamic computing via autonomous quantum thermal machines. In: Science advances, 2024, vol. 10, n° 36, p. eadm8792. doi: 10.1126/sciadv.adm8792

This publication URL: <https://archive-ouverte.unige.ch/unige:180590>

Publication DOI: [10.1126/sciadv.adm8792](https://doi.org/10.1126/sciadv.adm8792)

© The author(s). This work is licensed under a Creative Commons Attribution (CC BY 4.0)

<https://creativecommons.org/licenses/by/4.0>



COMPUTER SCIENCE

Thermodynamic computing via autonomous quantum thermal machines

Patryk Lipka-Bartosik*, Martí Perarnau-Llobet, Nicolas Brunner

We develop a physics-based model for classical computation based on autonomous quantum thermal machines. These machines consist of few interacting quantum bits (qubits) connected to several environments at different temperatures. Heat flows through the machine are here exploited for computing. The process starts by setting the temperatures of the environments according to the logical input. The machine evolves, eventually reaching a nonequilibrium steady state, from which the output of the computation can be determined via the temperature of an auxiliary finite-size reservoir. Such a machine, which we term a “thermodynamic neuron,” can implement any linearly separable function, and we discuss explicitly the cases of NOT, 3-MAJORITY, and NOR gates. In turn, we show that a network of thermodynamic neurons can perform any desired function. We discuss the close connection between our model and artificial neurons (perceptrons) and argue that our model provides an alternative physics-based analog implementation of neural networks, and more generally a platform for thermodynamic computing.

Copyright © 2024 the Authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original U.S. Government Works. Distributed under a Creative Commons Attribution License 4.0 (CC BY).

INTRODUCTION

Computing systems can take a variety of forms, from biological cells to massive supercomputers, and perform a broad range of tasks, from basic logic operation to machine learning. In all cases, the computational process must adhere to the principles of physics and, in particular, to the laws of thermodynamics. In general, information processing and thermodynamics are deeply connected, see, e.g., (1–3).

More recently, links between thermodynamics and computation are being developed. At the fundamental level, bounds for the thermodynamic cost of computation are derived, see, e.g., (4–6). From a more practical perspective, a promising direction explores low-dissipation computing. Here, models for elementary gates and circuits based on electronic transistors working in the mesoscopic regime, or even toward the single-electron mode, are considered (7–12). Crucially, thermodynamic models of computation must be thermodynamically consistent, meaning that they adhere to the laws of thermodynamics (13). This allows one to analyze their thermodynamic properties, e.g., energetic cost or dissipated heat, using the framework of stochastic thermodynamics. This approach already brought considerable progress, and further insight can be expected by moving to the fully quantum regime (14–18).

Another exciting direction is thermodynamic computing (19–22). This represents a paradigm for alternative physics-based models of computation, similarly to quantum computing or DNA computing. The main idea is to exploit the thermodynamic behavior of complex, nonequilibrium physical systems to perform computations, looking for not only a computational speedup but also a reduced energy cost. This approach has been explored in the context of machine learning and AI, see, e.g., (23–26). Very recently, promising progress has been reported, showing that a computational speedup in linear algebra problems can be achieved via a controllable system of coupled harmonic oscillators embedded in a thermal bath (27).

In this work, we develop a model for thermodynamic computing starting from a minimal model of a quantum thermal machine. More precisely, we develop autonomous quantum thermal machines that

can operate as computing devices where logical inputs and outputs are encoded in the temperature. As our device shares strong similarities with the basic model of an artificial neuron (the perceptron used, e.g., in neural networks), we refer to it as a “thermodynamic neuron.” Overall, our guiding motivation is to use diverse techniques offered by quantum thermodynamics to enhance our understanding of fundamental aspects of computation.

To construct our computing device, we start from the model of minimal autonomous quantum thermal machines (28, 29), which are made of a small quantum system (few interacting qubits) in contact with thermal baths at different temperatures. A first observation is that the effect of such a thermal machine onto an external system—heating or cooling—depends on the temperatures of the heat baths. Viewing these temperatures as an input and the temperature of the external system as an output, the thermal machine can be seen as a computing device (see Fig. 1). By associating a logical value to the temperature (e.g., cold temperatures corresponding to logical “0” and hot temperatures to logical “1”), we show that the autonomous machine can implement logical gates. As a first example, we show how a small quantum refrigerator/heat pump can be used to implement an inverter (NOT gate). This represents the simplest example of a thermodynamic neuron. In turn, we present a general model of a thermodynamic neuron and show that it can implement any Boolean linearly separable function. Such a function can be thought of as an assignment of 0 or 1 to the vertices of a Boolean hypercube (i.e., a geometric representation of its truth table). This allows one to divide the vertices into two sets. The Boolean function is said to be linearly separable if these two sets of points can be separated with a line. We discuss explicitly the examples of NOR and 3-MAJORITY. A key element in this construction is the concepts of virtual qubits and virtual temperatures (29), which allow us to establish a close connection between our machines and perceptrons, a common model of an artificial neuron. Furthermore, we show that, by constructing networks of thermodynamic neurons, one can implement any desired function, and we discuss the example of XOR. We detail an algorithm, inspired by artificial neural networks, for designing thermodynamic neurons (and their networks) for implementing any given target function. We conclude with a discussion and an outlook.

Department of Applied Physics, University of Geneva, 1211 Geneva, Switzerland.
*Corresponding author. Email: patryk.lipka.bartosik@gmail.com




Fig. 1. Thermodynamic neuron. The thermodynamic neuron is an autonomous quantum thermal machine designed for computing. The device consists of few interacting qubits (yellow dots), connected to several thermal environments. The input of the computation is encoded in the temperature of heat baths (depicted in red). This generates heat flows through the machine, which eventually reaches a nonequilibrium steady state. The output of the computation can be retrieved from the final temperature of a finite-size reservoir (shown in blue). By designing the machine (setting the qubit energies and their interaction), specific functions between the input and output temperatures can be implemented.

Before proceeding, we highlight a number of relevant features of our model. First, as it is constructed from a minimal model of quantum thermal machines, the model is thermodynamically consistent. Hence, the model allows for an examination of the trade-off between consumed energy, dissipation, and performance, which we investigate. Second, as it is based on changes of temperatures and flows of energy, the model involves only one conserved quantity, namely, energy. Computation in our model occurs solely as a result of heat flowing from one part of the machine to the other. This is in contrast to most conventional models of computation, in particular models for nanoscale electronic circuits, where heat is an unwanted by-product that hampers computation and introduces errors. Last, the functioning of our model can be intuitively understood by exploiting interesting connections between quantum systems at thermal equilibrium and artificial neural networks.

RESULTS

Framework

Autonomous quantum thermal machines

Quantum thermal machines usually consist of a small-scale physical system described within quantum theory. This system is then placed in contact with external resources, such as thermal baths or driving, to implement a thermodynamic task such as cooling, heating, or producing work; see, e.g., (30) or (31) for reviews.

Here, our focus is on a special class of quantum thermal machines known as autonomous quantum thermal machines [see (32) for a recent review]. Their main interest resides in the fact that these machines work autonomously, in the sense that they are powered by external resources that thermal (typically two or more heat baths at different temperatures) and their internal dynamics is time-independent (modeled via a time-independent Hamiltonian). While first models can be traced back to the thermodynamic analysis of masers (33), recent works have developed a framework for discussing minimal models of autonomous thermal machines, working as refrigerators, heat pumps, and heat engines (28, 29, 34). Many physical

models of quantum thermal machines (35–40) can be mapped back to these minimal abstract models (32). More recently, autonomous machines have also been devised for achieving other tasks such as the creation of entanglement (41), timekeeping (i.e., clocks) (42–44), and thermometry (45). A key aspect of these machines is their autonomy making them relevant from a practical perspective (46), and first proof-of-principle experiments have been reported (47, 48). More generally, the limits of designing autonomous quantum devices have been discussed (49).

Open quantum system dynamics

In this work, we will focus on autonomous quantum thermal machines consisting of few qubits, i.e., few two-level quantum systems. To start with, let us review the dynamics of a single qubit in contact with a heat bath. First, the qubit features two energy eigenstates: the ground state $|0\rangle$ and the excited state $|1\rangle$, with respective energies E_0 and $E_1 > E_0$. The state of the qubit is represented by a density operator ρ , and its mean energy is given by $\text{Tr}[\rho H]$, where $H = E_0|0\rangle\langle 0| + E_1|1\rangle\langle 1|$ denotes the Hamiltonian. A convenient quantity is the energy gap, $\epsilon := E_1 - E_0$. Without loss of generality we take $E_0 = 0$ so that the qubit's energy is fully specified by its energy gap. When placed in contact with an environment, the qubit evolution is described by the master equation

$$\dot{\rho} = -i[H, \rho] + \mathcal{L}[\rho] \quad (1)$$

The first term captures the unitary evolution governed by the Hamiltonian, while the second term captures the environment's impact on the qubit via the dissipator $\mathcal{L}[\cdot]$. Here, we use the common assumption of weak coupling to write down the dissipator, i.e., we assume that the qubit is weakly correlated with its environment.

As the qubit evolves over time, it eventually reaches a steady state when $\dot{\rho} = 0$. When the environment is a thermal bath, with an inverse temperature $\beta = 1/kT$, the resulting steady state is given by a qubit thermal (Gibbs) state: $\tau(\beta) = e^{-\beta H}/Z$, where $Z = \text{tr} e^{-\beta H}$ is the canonical partition function. In this case, the probability of the qubit to be in the excited state is given by the Fermi-Dirac distribution

$$g(\beta\epsilon) = \langle 1 | \tau(\beta) | 1 \rangle = \frac{1}{1 + e^{\beta\epsilon}} \quad (2)$$

Note that this function coincides with the sigmoid function used in machine learning. We will explore this connection more carefully later.

Thermal machines

The machines we will consider typically consist of several qubits with energy gaps ϵ_k . The qubits weakly interact with each other via an energy-preserving interaction. This is modeled by a time-independent interaction Hamiltonian, H_{int} , which commutes with the free Hamiltonian $H_0 = \sum_k \epsilon_k |1\rangle\langle 1|_k$, i.e., $[H_{\text{int}}, H_0] = 0$. In what follows, we will slightly abuse notation and write $|i\rangle\langle i|_k$ to denote a tensor product acting as identity everywhere except at position k , i.e., $1 \otimes \dots \otimes |i\rangle\langle i|_k \otimes \dots \otimes 1$. Each qubit is then connected to a thermal bath. In general, these baths are at different (inverse) temperatures β_k . When the coupling between qubits and thermal baths is weak, the dynamics of such a machine is well captured by a local master equation (50) of the form

$$\dot{\rho} = -i[H_0 + H_{\text{int}}, \rho] + \sum_k \mathcal{L}^{(k)}[\rho_k] \quad (3)$$

where ρ now denotes the multi-qubit state of the machine.

The main assumption that we are going to use in this work is local detailed balance, which, in our current context, means that local thermal states are the fixed point of each dissipator, i.e.

$$\mathcal{L}^{(k)}[\tau(\beta_k)] = 0 \quad (4)$$

This condition is well justified when the couplings in H_{int} are sufficiently weak (50). A quantity relevant to our analysis is the heat current released from the qubit to the heat bath in this process. This is given by

$$j_k := \text{Tr}\{H\mathcal{L}^{(k)}[\rho]\} \quad (5)$$

We note that, in certain cases, a qubit of the machine will be coupled to two different baths, in general, at different temperatures. In this case, the total dissipator for the qubit is simply obtained by summing the dissipators with respect to each bath. In turn, this implies that the total heat current is the sum of the heat currents with respect to each bath.

Although our key qualitative findings only require the detailed balance condition, introducing a specific thermalization model would allow us to support our results with numerical evidence. To keep the presentation simple, we will use the so-called reset model [see, e.g., (28)] in which the dissipators take the simple form $\mathcal{L}^{(k)}[\rho] = \gamma_k(\text{Tr}_k[\rho] \otimes \tau(\beta_k) - \rho)$, where $\text{Tr}_k[\cdot]$ denotes the partial trace over qubit k and γ_k is the coupling, which corresponds to the probability that qubit k thermalizes with its bath. We assume that all systems are labeled, and no relevance is given to the order of the tensor product. Note that $\text{Tr}_k[\rho] \otimes \tau_k(\beta_k)$ represents the multi-qubit state after a full thermalization event. This model can be viewed as a collisional process, where, in each instant of time, the qubit has a certain probability to collision with a thermal qubit from the bath. Within this model, the heat current from Eq. 6 is given by

$$j_k := \text{Tr}\{H\mathcal{L}^{(k)}[\rho]\} = \gamma_k \epsilon_k [g(\beta_k \epsilon_k) - p_k] \quad (6)$$

where p_k is the probability that the qubit connected to the bath is in an excited state. We note that, in certain cases, a qubit of the machine will be coupled to two different baths, in general, at different temperatures. In this case, the total dissipator for the qubit is simply obtained by summing the dissipators with respect to each bath. In turn, this implies that the total heat current is the sum of the heat currents with respect to each bath.

Last, a quantity of interest for our work is the dissipation generated by the machines. To quantify dissipation, we use entropy production rate $\dot{\Sigma}$. This quantity captures the fundamental irreversibility of the machine. The second law of thermodynamics restricts the behavior of any thermal machine. For our autonomous machines, it reads

$$\dot{\Sigma} := \dot{S}[\rho(t)] - \sum_k \beta_k j_k(t) \geq 0 \quad (7)$$

where $S(\rho) := -\text{Tr}[\rho \log \rho]$ is the von Neumann entropy of the machine and $j_k(t)$ is the total heat current flowing into the k th heat bath at time t . We also use the dot notation to indicate complete time derivatives, e.g., $\dot{\Sigma} \equiv \frac{d}{dt}\Sigma$.

The quantity $\dot{\Sigma}$ is the rate of entropy production, which quantifies the speed at which heat (entropy) is dumped into all environments connected with the machine, see, e.g., (30, 51, 52). It therefore

measures the amount of information that is lost (i.e., transferred to unobserved degrees of freedom). It is also a central quantity appearing in thermodynamic uncertainty relations (TURs) (53–55) as well as bounds on the speed of a stochastic evolution (56). We will be mostly interested in the dynamics of the steady state of the system, which corresponds to $\dot{\rho} = 0$ or equivalently $\dot{S}[\rho(t)] = 0$.

An important class of quantum thermal machines is autonomous machines. Such machines operate without requiring external control over their internal components (e.g., couplings or local energies) as they operate in the steady-state regime. This autonomy offers a key advantage: It eliminates the need for complex, high-precision control, which is a major contributor to the energy consumption of traditional nanoscale devices. An interesting platform for realizing autonomous quantum thermal machines are thermoelectric quantum dots (57).

Autonomous quantum thermal machines and virtual qubit

Before we explain our model of a computing thermal machine, it is worth discussing a simpler machine, namely, the three-qubit thermal machine introduced in refs. (28, 29). The intuition developed for this model can be then used to understand more complex quantum thermal machines.

Consider a thermal machine that consists of two qubits \mathcal{E}_0 and \mathcal{E}_1 such that \mathcal{E}_i is in a thermal contact with a heat bath at an inverse temperature β_i for $i = 0, 1$. Let ϵ_0 be the energy spacing of qubit \mathcal{E}_0 and $\epsilon_1 \leq \epsilon_0$ be the energy spacing of qubit \mathcal{E}_1 . In the absence of interactions with an external system, each qubit interacts only with its own thermal bath and hence reaches thermal equilibrium at the corresponding inverse temperature. Therefore, the state of qubit \mathcal{E}_i can be written as

$$\tau_{\mathcal{E}_i}(\beta_i) = \frac{1}{Z_{\mathcal{E}_i}} (|0\rangle\langle 0|_{\mathcal{E}_i} + e^{-\beta_i \epsilon_i} |1\rangle\langle 1|_{\mathcal{E}_i}) \quad (8)$$

where $Z_{\mathcal{E}_i} = 1 + e^{-\beta_i \epsilon_i}$. Consequently, the two qubits are jointly described by a tensor product state $\tau_{\mathcal{E}_0}(\beta_0) \otimes \tau_{\mathcal{E}_1}(\beta_1)$ and have four different energy eigenstates, i.e., $|i\rangle_{\mathcal{E}_0} |j\rangle_{\mathcal{E}_1}$ for $i, j \in \{0, 1\}$. Let us now focus on two particular eigenstates, namely

$$|0\rangle_v := |0\rangle_{\mathcal{E}_0} |1\rangle_{\mathcal{E}_1}, \quad |1\rangle_v := |1\rangle_{\mathcal{E}_0} |0\rangle_{\mathcal{E}_1} \quad (9)$$

These two states have an energy spacing $\epsilon_v := \epsilon_1 - \epsilon_0$ and span a subspace of the joint Hilbert space that is usually referred to as the virtual qubit [see also (58)]. For that subspace, we can further assign a virtual temperature β_v by looking at the ratio of populations in the virtual qubit, that is

$$e^{-\beta_v \epsilon_v} := \frac{v \langle 1 | \tau_{\mathcal{E}_0}(\beta_0) \otimes \tau_{\mathcal{E}_1}(\beta_1) | 1 \rangle_v}{v \langle 0 | \tau_{\mathcal{E}_0}(\beta_0) \otimes \tau_{\mathcal{E}_1}(\beta_1) | 0 \rangle_v} = \frac{e^{-\beta_0 \epsilon_0}}{e^{-\beta_1 \epsilon_1}} \quad (10)$$

which allows us to express β_v as

$$\beta_v = \left(\frac{\epsilon_0}{\epsilon_0 - \epsilon_1} \right) \beta_0 - \left(\frac{\epsilon_1}{\epsilon_0 - \epsilon_1} \right) \beta_1 \quad (11)$$

Observe that the virtual temperature, as a function of the local energies ϵ_0 and ϵ_1 , can take any range of values. In particular, notice that it can fall outside of the range specified by β_0 and β_1 and can even take negative values. This corresponds to a population inversion (29).

Let us now add another qubit to the machine, denote it with \mathcal{E}_2 and place it in a thermal contact with the virtual qubit. To enable an interaction between the new qubit and the virtual qubit, we choose

the energy of the former to be $\epsilon_z = \epsilon_v$. This allows the systems to resonantly exchange energy with the following Hamiltonian

$$H_{\text{int}} = \chi |1\rangle\langle 0|_v \otimes |0\rangle\langle 1|_{\mathcal{E}_z} + \text{h.c.} \quad (12)$$

where χ specifies the coupling strength and “h.c.” stands for Hermitian conjugate. The above interaction induces a transition between two degenerate energy eigenstates $|1\rangle_v|0\rangle_{\mathcal{E}_z} \leftrightarrow |0\rangle_v|1\rangle_{\mathcal{E}_z}$, which effectively places the virtual qubit in a thermal contact with the new qubit. After a sufficiently long amount of time, the temperature of the qubit \mathcal{E}_z reaches the virtual temperature β_v .

Let us now observe that such a three-qubit thermal machine can operate as either a refrigerator or a heat pump. In Fig. 2, we plot the virtual (inverse) temperature β_v as a function of β_1 for a fixed value of β_0 . More specifically, notice that, when $\beta_1 \leq \beta_0$, the inverse virtual temperature is larger than both β_0 and β_1 ; hence, the machine operates as a refrigerator. When $\beta_0 < \beta_1 < (\epsilon_0/\epsilon_1)\beta_0$, the inverse virtual temperature β_v is smaller than both β_0 and β_1 ; hence, the machine is a heat pump. Last, when $\beta_1 > (\epsilon_0/\epsilon_1)\beta_0$, the virtual temperature is negative, meaning that the device operates as a heat engine. Notice that, in all three regimes, the virtual temperature falls outside of the range of “easily accessible” temperatures specified by β_0 and β_1 that could be achieved simply by coupling one of the qubits to the two heat baths.

Thermodynamic neuron for NOT gate

In this section, we describe an autonomous thermal machine implementing an inverter (NOT gate). This represents the simplest example of a thermodynamic neuron. We start with a short and intuitive description of the machine’s operation after which we provide a more in-depth discussion of its functioning.

The machine is sketched in Fig. 3A. It is composed of two parts, which we refer to as the collector (\mathcal{C}) and the modulator (\mathcal{M}). The collector consists of three interacting qubits connected to different environments. The first two qubits (denoted \mathcal{E}_0 and \mathcal{E}_1) are connected to two heat baths, denoted \mathcal{B}_0 and \mathcal{B}_1 , at inverse temperatures β_0 and β_1 , respectively. The first bath \mathcal{B}_0 simply represents a reference bath; hence, β_0 will simply be fixed to a certain value and called the reference temperature. The second bath \mathcal{B}_1 will be used to

encode the input of the computation. These two heat baths are supposed to have an infinitely large heat capacity; hence, their temperature will remain constant during the time evolution of the machine. Last, the third qubit of the collector (denoted \mathcal{E}_z) is connected to an environment \mathcal{B}_z with a finite heat capacity C (this can be viewed as a finite-size reservoir). The key point is that the inverse temperature β_z of \mathcal{B}_z will evolve in time, and the final temperature (in the steady-state regime) will encode the output of the computation.

To guide intuition, it is useful to think of the collector as a simple (three-qubit) thermal machine (28, 29), which we discussed in Autonomous quantum thermal machines and virtual qubit. When the input temperature is hot ($\beta_1 < \beta_0$), the machine works as a refrigerator, i.e., cooling down the output environment \mathcal{B}_z . On the contrary, when the input temperature is cold ($\beta_1 > \beta_0$), the machine works as a heat pump, heating up \mathcal{B}_z . Hence, we see that the machine works as a sort of inverter for the temperature. We encourage the reader to take a look at Fig. 2, which illustrates different regimes of operation of a three-qubit machine that is equivalent to the collector of the NOT gate.

Because of the action of the collector, the output inverse temperature β_z depends linearly on the input β_1 , as demonstrated in Eq. 11. From a signal processing perspective, this translates to an inverting linear amplifier. When a signal passes through a sequence of such devices, any noise present in the signal will be amplified, potentially leading to unwanted bit flips. To enhance the noise robustness of the collector, a nonlinear modulation of the output inverse virtual temperature is required. This modulation should minimize the output variation for small input fluctuations within designated logical regions (i.e., where the collector acts as a refrigerator or a heat pump). At the same time, the output should change substantially when the input transitions to a different logical region. This ensures that the any noise-induced distortion of the signal in the output will be minimized.

The above modulation will be realized by another part of the machine, i.e., the modulator. It is a single qubit machine that is coupled to two thermal baths, i.e., a reference bath \mathcal{B}_r with a fixed inverse temperature β_r and the output bath \mathcal{B}_z (see Fig. 3A). This has the effect to delimit a specific range for the output temperatures β_z , making the response of the device effectively nonlinear and hence closer to an ideal NOT gate.

In the following, we present in detail the models for the collector and the modulator and then discuss the dynamics of the machine and its operation as a NOT gate. Last, we investigate the trade-off between the gate performance (as given by the average error rate) and dissipation (as given by entropy production).

Collector

The collector \mathcal{C} is composed of three qubits, which we denote \mathcal{E}_i for $i \in \{0, 1, z\}$ (see Fig. 4), with energy gaps ϵ_i . Each qubit is weakly coupled to an environment, denoted \mathcal{B}_i , at (inverse) temperatures β_i with the coupling constants γ for \mathcal{E}_0 and \mathcal{E}_1 and μ for \mathcal{E}_z . Therefore, the collector can be seen as a three-qubit thermal machine that we discussed in Autonomous quantum thermal machines and virtual qubit. This three-qubit system is described by a joint state $\rho_{\mathcal{C}}$ that evolves according to the master equation (Eq. 3), i.e.

$$\dot{\rho}_{\mathcal{C}} = -i[H_0 + H_{\text{int}}, \rho_{\mathcal{C}}] + \mathcal{L}[\rho_{\mathcal{C}}] \quad (13)$$

with the local Hamiltonian $H_0 = \sum_{i \in \{0, 1, z\}} \epsilon_i |1\rangle\langle 1|_{\mathcal{E}_i}$ and local dissipators $\mathcal{L} = \mathcal{L}^{(0)} + \mathcal{L}^{(1)} + \mathcal{L}^{(z)}$.

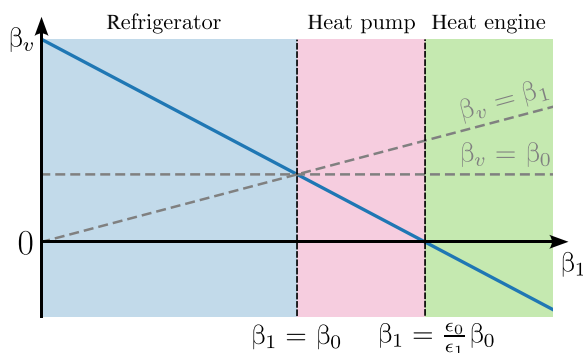


Fig. 2. Different operation regimes of a three-qubit thermal machine. The plot shows the inverse virtual temperature β_v as a function of bath inverse temperature β_1 when keeping β_0 fixed. When $\beta_1 < \beta_0$, the inverse virtual temperature becomes larger than both β_0 and β_1 , which means that the machine operates as a refrigerator. When $\beta_0 < \beta_1 < (\epsilon_0/\epsilon_1)\beta_0$, we have the exactly opposite situation and the machine operates as a heat pump. Last, when $\beta_1 > (\epsilon_0/\epsilon_1)\beta_0$, the machine operates as a heat engine. Figure adapted from ref. (29).

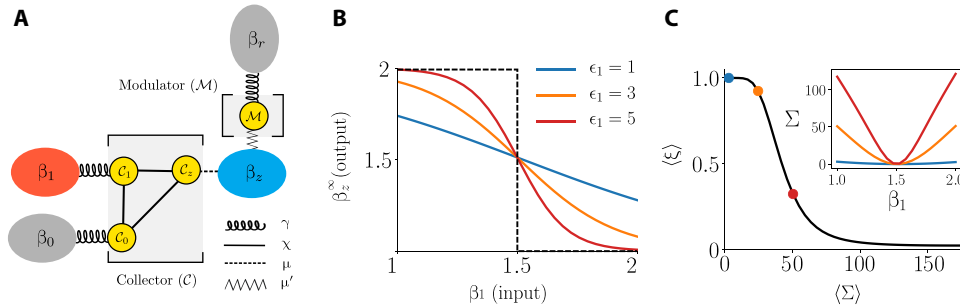


Fig. 3. Thermodynamic neuron for implementing a NOT gate. (A) Design of the machine. The collector consists of three interacting qubits (yellow dots), each connected to a thermal environment. The logical input is encoded in the temperature β_1 of the heat bath \mathcal{B}_1 (red), while the output will be retrieved from the final temperature β_z^∞ of the finite-size reservoir \mathcal{B}_z (blue); the heat bath \mathcal{B}_0 is at a fixed reference temperature. The collector implements the desired inversion of the temperature. To make the response nonlinear, we must add the modulator, which consists of an additional qubit connected to a reference heat bath. (B) Relation between the input temperature β_1 and the final output temperature β_z^∞ (in the steady-state regime). Notably, the machine produces the desired inversion of the temperature. The quality of the response can be increased by tuning the machine parameters, in particular by increasing the energy gap ϵ_1 of the collector qubit C_1 . Black dashed line shows the characteristics of an ideal NOT gate. (C) Trade-off between the average dissipation $\langle \Sigma \rangle$ (see Eq. 29) and the average error $\langle \xi \rangle$ (see Eq. 27). We see clearly that, to increase robustness to noise, the machine must dissipate more heat to the environment. The inset shows the entropy production as a function of the input temperature β_1 for different values of the qubit energy ϵ_1 . Parameter values: $\beta_{\text{hot}} = 1$, $\beta_{\text{cold}} = 2$, $\gamma = \chi = 1$, $\mu = 10^{-4}$, $\epsilon_z = 0.1$, $\tau = 10^8$, and $\beta_0 = \beta_z(0) = 3/2$.




Fig. 4. Virtual qubit in the collector. The sketch shows the energy structure of a three-qubit collector. The Hilbert space of the two physical qubits C_0 and C_1 contains a two-dimensional subspace with an energy gap $\epsilon_v = \epsilon_z$ (so-called virtual qubit) and effective temperature β_v (so-called virtual temperature). The interaction Hamiltonian H_{int} is chosen so that this virtual qubit interacts with the physical qubit C_z , cooling it down (or heating up) in the process.

It is important to ensure that energy can flow between the qubits. For this, we choose the energy gap of the third qubit \mathcal{C}_z to be $\epsilon_z = \epsilon_0 - \epsilon_1$. This implies that the two states $|1\rangle_{\mathcal{C}_0}|0\rangle_{\mathcal{C}_1}|0\rangle_{\mathcal{C}_z}$ and $|0\rangle_{\mathcal{C}_0}|1\rangle_{\mathcal{C}_1}|1\rangle_{\mathcal{C}_z}$ have the same energy and can be coupled via the interaction Hamiltonian

$$H_{\text{int}} = \chi |1\rangle_{\mathcal{C}_0} \langle 0|_{\mathcal{C}_0} \otimes |0\rangle_{\mathcal{C}_1} \langle 1|_{\mathcal{C}_1} \otimes |0\rangle_{\mathcal{C}_z} \langle 1|_{\mathcal{C}_z} + \text{h.c.} \quad (14)$$

where χ is the coupling strength. This interaction conserves the total energy (because $[H_0, H_{\text{int}}] = 0$), which guarantees that energy can be exchanged even in the weak coupling regime.

We want to understand the effect of the collector on the output environment \mathcal{B}_z in the steady-state regime, i.e., when $\dot{\rho}_{\mathcal{B}_z} = 0$. To do so, we will follow the approach of ref. (29), which is summarized in Autonomous quantum thermal machines and virtual qubit and visualized in Fig. 4.

First, note that, from the form of the interaction Hamiltonian H_{int} , we see that there are only two states of the machine that exchange energy in the steady-state dynamics. These are simply the two states we discussed above that have the same energy. Now let us think of the three-qubit system as a machine comprising the first two qubits \mathcal{C}_0 and \mathcal{C}_1 and the target qubit \mathcal{C}_z . The effect of the machine is to thermalize the target qubit \mathcal{C}_z with a virtual qubit characterized by the two levels

$$|0\rangle_v := |0\rangle_{\mathcal{C}_0}|1\rangle_{\mathcal{C}_1}, \quad |1\rangle_v := |1\rangle_{\mathcal{C}_0}|0\rangle_{\mathcal{C}_1} \quad (15)$$

These levels form a virtual qubit with energy gap $\epsilon_v = \epsilon_0 - \epsilon_1$. Let us denote with $g_v := \langle 1|_v \tau_{\mathcal{C}_0}(\beta_0) \otimes \tau_{\mathcal{C}_1}(\beta_1) |1\rangle_v$ the occupation of the excited state of this effective system. Then, the ratio of populations in the subspace associated with the virtual qubit becomes $g_v/(1 - g_v) = e^{-\beta_v(\epsilon_0 - \epsilon_1)}$, where β_v is the (inverse) virtual temperature

$$\beta_v = \left(\frac{\epsilon_0}{\epsilon_0 - \epsilon_1} \right) \beta_0 - \left(\frac{\epsilon_1}{\epsilon_0 - \epsilon_1} \right) \beta_1 \quad (16)$$

Using the intuition developed in Autonomous quantum thermal machines and virtual qubit, we can now understand the steady-state dynamics of the collector \mathcal{C} . The collector aims to thermalize the target qubit \mathcal{C}_z to the virtual temperature β_v , as can be seen by re-writing the interaction Hamiltonian in terms of virtual qubit levels as in Eq. 12. The only difference with respect to the setting from Autonomous quantum thermal machines and virtual qubit is that now the target qubit is itself coupled to a finite heat bath \mathcal{B}_z at an inverse temperature β_z and therefore the regime of the collector's operation (i.e., if it acts as a refrigerator or a heat pump) is defined with respect to the inverse temperature β_z instead of β_0 .

When $\beta_v > \beta_z$, energy flows from the target qubit \mathcal{C}_z to the machine (via the virtual qubit), effectively cooling the target qubit down; the machine acts as a refrigerator. On the other hand, when $\beta_v < \beta_z$, energy flows toward the qubit \mathcal{C}_z , heating it up in the process;

the machine acts as a heat pump. Which one of these different machine's behaviors actually occurs depending on the inverse temperatures β_0 and β_1 via Eq. 16? This ability of the collector to change its behavior based on the input temperature is the basic principle behind our inverter.

Recall that the target qubit \mathcal{C}_z is coupled to its own (finite) thermal bath \mathcal{B}_z . In turn, the mechanism described above will have the effect of thermalizing the output environment \mathcal{B}_z to the virtual temperature. To see this, consider the steady-state current from the collector \mathcal{C} to the output environment \mathcal{B}_z under the reset model of thermalization (see Eq. 6)

$$j_{\mathcal{C}} := \mu \epsilon_z [g_z(\beta_z) - g_z(\beta_v)] \quad (17)$$

where $g_z(x) := g(x\epsilon_z)$ and g is the Fermi-Dirac distribution from Eq. 2. The collector attempts to bring the temperature of the environment \mathcal{B}_z closer to the virtual temperature. By choosing energy gaps ϵ_0 and ϵ_1 appropriately [i.e., the linear weights in Eq. 16], we can, in principle, obtain any linear inverting behavior.

Modulator

The modulator \mathcal{M} is composed of a single qubit with an energy gap $\epsilon_{\mathcal{M}} = \epsilon_z$. The qubit is put in contact with two thermal baths: \mathcal{B}_r at an inverse temperature β_r with a coupling rate γ and \mathcal{B}_z with a different coupling rate μ' . The qubit state $\rho_{\mathcal{M}}$ evolves according to the following master equation

$$\dot{\rho}_{\mathcal{M}} = \mathcal{L}^{(r)}[\rho_{\mathcal{M}}] + \mathcal{L}^{(z)}[\rho_{\mathcal{M}}] \quad (18)$$

In the steady state, the excited-state population of the qubit depends only on the coupling rates γ and μ' . We set these rates so that $\mu' \ll \gamma$, ensuring that the qubit will effectively thermalize to the inverse temperature β_r . Therefore, the steady-state heat current from \mathcal{M} to \mathcal{B}_z under the reset model (Eq. 6) reads

$$j_{\mathcal{M}} := \mu' \epsilon_z [g_z(\beta_z) - g_z(\beta_r)] \quad (19)$$

The modulator attempts to bring β_z closer to the (inverse) temperature β_r , and the strength of this effects is controlled by the coupling rate μ' . The choice of the values of β_r and μ' will therefore completely specify the behavior of the modulator. By appropriately choosing these two parameters, we can specify the range of the output temperature β_z leading to a nonlinear response of the machine (see the Supplementary Materials A for more details).

Dynamics of the machine

We now combine our understanding of the collector and the modulator to gain insight into the full evolution of the machine. The collector and the modulator are both connected to an environment \mathcal{B}_z with a finite heat capacity C . The temperature change of this environment is proportional to the sum of all entering heat currents. Specifically, we assume that the temperature $T_z := 1/\beta_z$ changes according to the calorimetric equation $\dot{T}_z = \frac{1}{C}(j_{\mathcal{C}} + j_{\mathcal{M}})$, which, in terms of β_z , reads

$$\dot{\beta}_z = -\frac{1}{C}\beta_z^2(j_{\mathcal{C}} + j_{\mathcal{M}}) \quad (20)$$

Consequently, the steady-state inverse temperature β_z^∞ is obtained by solving the equation $j_{\mathcal{C}} + j_{\mathcal{M}} = 0$.

Crucially, the couplings of the collector and the modulator to \mathcal{B}_z are set to be much weaker than their couplings to the heat baths \mathcal{B}_0 , \mathcal{B}_1 , and \mathcal{B}_r , i.e., we have that $\gamma \gg \mu, \mu'$. This implies that the

dynamics of the whole machine has two intrinsic timescales. The first (fast dynamics) is associated with the internal evolution of the collector \mathcal{C} and the modulator \mathcal{M} . Hence, both parts of the machine will reach their steady states relatively quickly. This means that the qubit \mathcal{C}_z of the collector will reach the virtual temperature β_v (see Eq. 16), while the modulator qubit will be at temperature β_r . The second (slow dynamics) is associated with the changes of the temperature of the output environment \mathcal{B}_z . This means that \mathcal{B}_z will slowly thermalize via the contact with qubits \mathcal{C}_z and \mathcal{M} to an intermediate temperature between β_v and β_r .

Let us now discuss the slow evolution more carefully. We denote by $\beta_z(t)$ the time evolution of the temperature of the output environment \mathcal{B}_z . The heat currents delivered from the collector and the modulator alter $\beta_z(t)$ according to Eq. 20. The steady state of the output environment \mathcal{B}_z is achieved when $\dot{\beta}_z(t) = 0$. Denoting the stationary value of $\beta_z(t)$ with β_z^∞ and solving the equation $j_{\mathcal{C}} + j_{\mathcal{M}} = 0$, we obtain the following expression for the steady-state temperature

$$g_z(\beta_z^\infty) = \Delta g_z(\beta_v) + (1 - \Delta)g_z(\beta_r) \quad (21)$$

where $\Delta := \mu/(\mu' + \mu)$. To interpret the temperature of the output reservoir \mathcal{B}_z as a valid logical signal, we need to limit the possible values of output temperature β_z^∞ to a well-defined range β_{cold} and β_{hot} , where the parameters satisfy $\beta_{\text{cold}} > \beta_{\text{hot}}$ but are otherwise arbitrary. To enforce this requirement, we can fix the free parameters of the modulator (see the Supplementary Materials A for details). Choosing μ' and β_r so that $\Delta = g_z(\beta_{\text{hot}}) - g_z(\beta_{\text{cold}})$ and $g_z(\beta_r) = g_z(\beta_{\text{cold}})/(1 - \Delta)$ leads to

$$\beta_z^\infty = \frac{1}{\epsilon_z} \log \left[Q(\beta_v)^{-1} - 1 \right] \quad (22)$$

with $Q(\beta_v) := g_z(\beta_{\text{hot}})g_z(\beta_v) + g_z(\beta_{\text{cold}})[1 - g_z(\beta_v)]$ and β_v is the virtual temperature given in Eq. 16.

At this point, we are ready to discuss the performance of our inverter. In Fig. 3B, we plot the transfer characteristics (TC) of our machine in the steady-state regime. Specifically, we see that the behavior between the input and the output temperatures, β_1 and β_z , respectively, is an inversion. For a cold (hot) input temperature, the output temperature is hot (cold). Note that, in the figure, we have set $\beta_{\text{cold}} = 2$ and $\beta_{\text{hot}} = 1$. More generally, from Eq. 22, we see that, (i) when $\beta_1 = \beta_{\text{hot}}$, we have $\beta_z^\infty \approx \beta_{\text{cold}}$, and (ii) when $\beta_1 = \beta_{\text{cold}}$, we get $\beta_z^\infty \approx \beta_{\text{hot}}$.

In addition, we can see from the figure that the quality of the NOT gate depends on the model parameters, in particular on the energy gap ϵ_1 of the collector qubit \mathcal{C}_1 . The larger ϵ_1 becomes, the closer we get to an ideal NOT gate (i.e., inverted step function). It can be shown that, in the limit $\epsilon_1 \rightarrow \infty$, the TC becomes the ideal inverted step function. We investigate analytically in the Supplementary Materials A the properties of the TC in Eq. 22, showing its dependence on the energies of the collector qubits ϵ_0 and ϵ_1 and the inverse temperature β_0 of the reference bath. More specifically, Eq. 22 describes a function that is very similar to a sigmoid (or Fermi-Dirac) function $f(x) = (1 + e^x)^{-1}$, i.e.

$$\beta_z^\infty = f(x) + \mathcal{O}(\epsilon_z) \quad (23)$$

where $x := (\epsilon_1 + \epsilon_z)(\beta_0 - \beta_1)$. When ϵ_z is small (compared to ϵ_1), the roles of the free parameters become clear: β_0 characterizes the location of the step in β_z^∞ and $\epsilon_0 \approx \epsilon_1$ describes its steepness. For larger values of ϵ_z , the TC still demonstrates the desired inverting behavior; however, the role of the parameters ϵ_0 and β_0 becomes a bit more

complicated to interpret (see the Appendix Supplementary Materials A for details).

We note that the exact functional dependence between the input β_1 and the output β_z^∞ depends on the amount of heat current and hence also on the explicit thermalization model used. To arrive at Eq. 23, we used the simple reset model from Eq. 6. Choosing a different thermalization model leads to different mathematical forms (nonlinear functions f); however, the machine's fundamental ability to invert temperatures remains unchanged.

Logic operation

As seen above, our device produces the desired inversion relation between the input and output temperatures. The next step is to use the machine as a NOT gate, for which we must now encode the logical information appropriately in the corresponding temperatures.

In what follows, the input and output signals will be described by random variables $x, y \in \{0, 1, \emptyset\}$, where 0, 1 represent the binary logical values and \emptyset denotes an invalid result that cannot be assigned. The logical input x is encoded in the inverse temperature β_1 of heat bath \mathcal{B}_1 , while the logical output y is decoded from the final (inverse) temperature β_z^∞ of \mathcal{B}_z . For that, we use the mapping

$$x = \begin{cases} 0, \beta_1 = \beta_{\text{hot}} \\ 1, \beta_1 = \beta_{\text{cold}} \\ \emptyset \text{ otherwise} \end{cases} \quad y = \begin{cases} 0, \beta_z^\infty \leq (1 + \delta)\beta_{\text{hot}} \\ 1, \beta_z^\infty \geq (1 - \delta)\beta_{\text{cold}} \\ \emptyset \text{ otherwise} \end{cases} \quad (24)$$

Parameters β_{cold} and β_{hot} characterize the machine's range of operation, while δ captures its robustness to noise in the output signal. All these parameters are a part of the machine's design and can be chosen arbitrarily, depending on the specific working conditions (e.g., how much noise is the machine expected to tolerate). Mapping logical values to intervals as above allows one to tolerate fluctuations in the output signal, i.e., interpret them correctly even if they differ between rounds due to the stochasticity of the machine's evolution. In principle, we could also consider having noise in the input signal. Similarly, we could also consider mapping the output of the machine to several logic states, therefore effectively simulating a function with several output values. However, to keep the presentation simple, we will not do this here.

The thermal machine discussed in Thermodynamic neuron for NOT gate performs computation in an inherently stochastic manner, and therefore the actual machine's output will fluctuate around the steady-state value from Eq. 22. This will lead to possible errors in the gate implementation. Characterizing these errors is important to assess the quality of the gate, in terms of its robustness to noise.

In the following, we describe the machine as a binary channel defined by the encoding $e(\beta_1 | x)$ and decoding $d(y | \beta_z^\infty)$ as specified in Eq. 24. The input distribution is denoted $p(x)$. The behavior of the machine is then specified by a conditional distribution

$$p(y | x) := \int d(y | \beta_z) T(\beta_z | \beta_1) e(\beta_1 | x) d\beta_1 d\beta_z \quad (25)$$

where $T(\beta_z | \beta_1)$ describes the actual response β_z of the machine to the input β_1 . Because the evolution is ultimately stochastic, we assume that the actual response of the machine to the input β_1 is distributed according to

$$T(\beta_z | \beta_1) \propto \mathcal{N}(\beta_z^\infty, C) \quad (26)$$

where $\mathcal{N}(\mu, \sigma)$ is a Gaussian with mean μ and SD σ . The output heat bath \mathcal{B}_z is a macroscopic system that is composed of a large number

of particles. In such a large system, according to the central limit theorem, the sum of temperature fluctuations tends toward a Gaussian distribution. Because the temperature β_z is a macroscopic property related to the average kinetic energy of the particles, it reflects the sum of these microscopic fluctuations, and hence Gaussian distribution provides a reasonable approximation to the actual distribution. Moreover, a larger heat bath (higher C) can sustain more energy fluctuations without a substantial change in its average temperature. This translates to a wider distribution (larger SD) in the Gaussian distribution.

The average computation error $\langle \xi \rangle$ is the probability of observing an output different from the desired one, i.e.

$$\langle \xi \rangle = \sum_{x \in \{0, 1\}} \sum_{y \in \{0, 1\}} p(x) p(y | x) \delta_{xy} \quad (27)$$

where δ_{xy} is the Kronecker delta function. The above quantity is directly related to the shape of the TC (see Fig. 3B). Notably, the closer TC is to an ideal NOT gate (black dashed line), the smaller is $\langle \xi \rangle$. The actual TC of our machine approaches the ideal one in the limit of $\epsilon_1 \rightarrow \infty$. This indicates that the quality of the computation can be enhanced at the cost of using more energy, which implies that the machine will dissipate more heat. In the following discussion, we will examine this trade-off in more detail.

Trade-off between entropy production and noise robustness

Here, we investigate the relation between the quality of the gate, as quantified by the average computation error, to its thermodynamic cost, given by the amount of entropy that is produced during the computation.

First, let us evaluate the entropy production. As mentioned, the dynamics of the machine features two different timescales. The primary source of dissipation is given by the slow dynamics, in which the temperature of the output reservoir changes. The latter being connected to the collector and the modulator, the total dissipation rate is given by $\dot{\Sigma} = \dot{\Sigma}_{\mathcal{C}} + \dot{\Sigma}_{\mathcal{M}}$. We have that $\dot{\Sigma}_{\mathcal{C}} = -\beta_0 j_0 - \beta_1 j_1 - \beta_z j_{\mathcal{C}}$ and $\dot{\Sigma}_{\mathcal{M}} = -\beta_z j_{\mathcal{M}} - \beta_r j_r$; here, $j_{0,1,r}$ denotes the current from the heat bath $\mathcal{B}_{0,1,r}$ to their respective qubit. Under the action of the slow dynamics, the entropy of the qubits in the machine does not change, i.e., $\dot{S}(\rho_S) = 0$. Because of this, the entropy production is the weighted sum of the heat dissipated in each environment. To quantify the total dissipation incurred during the computation, we have to integrate the dissipation rate over time, i.e.

$$\langle \Sigma(\beta_1) \rangle = \int_0^\tau \langle \dot{\Sigma}(\beta_1) \rangle dt \quad (28)$$

where τ is the running time of the computation, indicating when the final temperature output β_z^∞ is read off.

We see that this quantity depends on β_1 . Hence, the dissipation will vary depending on the input. In the inset of Fig. 3C, we show this behavior, also considering different values of the parameter ϵ_1 . As expected, because the rate of dissipation is proportional to the heat currents flowing into the environments, the larger the energy of the qubits, the larger the rate of heat dissipation. Moreover, as expected, when $\beta_1 = \beta_0$, dissipation vanishes.

Next, let us estimate the dissipation averaged over different rounds of the computation, i.e., averaging over the inputs. We get the quantity

$$\langle \Sigma \rangle = \sum_{x \in \{0, 1\}} \int p(x) e(\beta_1 | x) \Sigma(\beta_1) d\beta_1 \quad (29)$$

In Fig. 3C, we examine the relation between the total dissipation $\langle \Sigma \rangle$ and the average computation error $\langle \xi \rangle$. We consider a uniform input distribution $p(x) = 1/2$ and a sufficiently long computing time to ensure we are close to the steady-state regime, $\tau = 10^8$. There is a monotonous relation between the two quantities. As expected, we see that lowering the average error rate comes at the price of increasing the dissipation.

Last, let us mention that the choice of temperatures β_{cold} and β_{hot} plays a crucial role in the machine's performance. A wider temperature range enhances noise resistance, making the machine less susceptible to small temperature fluctuations. However, this benefit comes at a cost—a larger temperature gap increases the thermalization time, essentially slowing down the thermodynamic neuron's computations. This creates an interesting trade-off between noise robustness and computational speed.

Thermodynamic neuron for linearly separable functions

In the previous section, we presented an autonomous thermal machine for performing a simple computation task, namely, inverting a signal. In this section, we generalize this construction for performing more complex computations. In particular, we show that any linearly separable function (from n bits to one bit) can be implemented via such a machine and give an effective algorithm for setting the appropriate machine parameters. This represents the general form of a thermodynamic neuron. We discuss explicitly examples for implementing the NOR gate and 3-MAJORITY.

A key step will be to establish a close connection between the thermodynamic neuron and the perceptron, the standard algorithm for modeling an artificial neuron. In particular, this connection exploits the notion of the virtual qubit.

Model

A thermodynamic neuron is an autonomous quantum thermal machine that implements a binary function from n bits to one bit. In analogy with the thermal machine for inversion from

Thermodynamic neuron for NOT gate, the general model of a thermodynamic neuron consists of two main parts: the collector \mathcal{C} and the modulator \mathcal{M} (see Fig. 5). The design of the collector is a generalization of the single-input collector, while the modulator is exactly the same.

The (generalized) collector \mathcal{C} consists of $n + 2$ qubits \mathcal{C}_i with energy gaps ϵ_i . The first qubit \mathcal{C}_0 is connected to the reference heat bath \mathcal{B}_0 at a fixed inverse temperature β_0 . The remaining qubits \mathcal{C}_1 to \mathcal{C}_n are connected to input heat baths, their temperatures (β_1 to β_n) encoding the n input bits. The last qubit \mathcal{C}_z is connected to the output reservoir \mathcal{B}_z with a finite heat capacity C . The modulator \mathcal{M} consists of a single qubit connected to a heat bath at a reference temperature β_r and the output reservoir \mathcal{B}_z .

To understand the dynamics of the collector, we will again use the idea of a virtual qubit, now associated with a two-dimensional subspace within the Hilbert space of qubits $\mathcal{C}_0, \dots, \mathcal{C}_n$. A multi-qubit machine can have many virtual qubits; hence, we need notation to specify which virtual qubit is relevant for our problem. For that, we introduce a binary vector $\mathbf{h} = (h_0, h_1, \dots, h_n)$, where $h_i \in \{0, 1\}$ denotes if a given physical qubit i contributes its ground ($h_i = 0$) or excited ($h_i = 1$) state to the virtual qubit.

A virtual qubit specified by a vector \mathbf{h} consists of two multi-qubit energy levels

$$|0\rangle_v := |h_0\rangle_{\mathcal{C}_0} |h_1\rangle_{\mathcal{C}_1} \dots |h_n\rangle_{\mathcal{C}_n} \quad (30)$$

$$|1\rangle_v := |h_0 \oplus 1\rangle_{\mathcal{C}_0} |h_1 \oplus 1\rangle_{\mathcal{C}_1} \dots |h_n \oplus 1\rangle_{\mathcal{C}_n} \quad (31)$$

where \oplus denotes addition mod 2. The energy gap ϵ_v of the virtual qubit with levels $|0\rangle_v$ and $|1\rangle_v$ is given by

$$\epsilon_v := \sum_{i=0}^n (-1)^{h_i \oplus 1} \epsilon_i \quad (32)$$

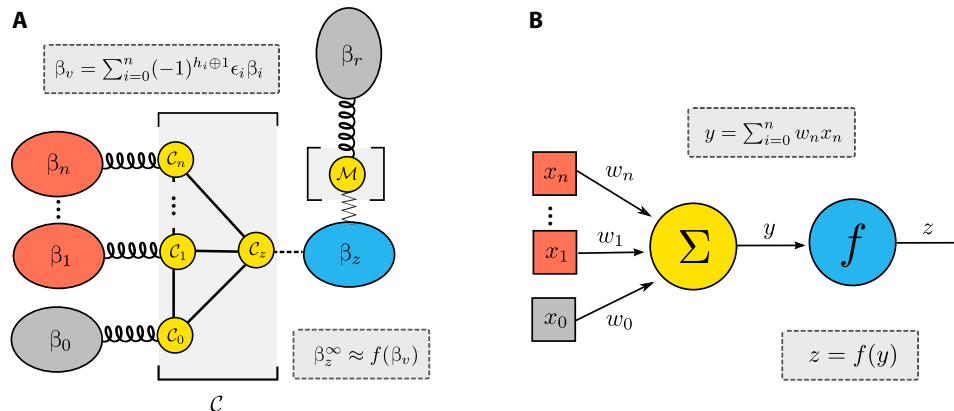


Fig. 5. General model of the thermodynamic neuron and analogy with a perceptron. (A) Structure of a thermodynamic neuron for implementing an n -to-one bit function. The collector \mathcal{C} consists of $n + 2$ qubits, connected to the input heat baths (red), reference heat baths (gray), and the output reservoir (blue). The working principle of the collector is to thermalize qubit \mathcal{C}_z to the virtual temperature β_v (see Eq. 29). In turn, this affects the temperature of the finite-size output reservoir \mathcal{B}_z (blue). The modulator controls the range of output temperatures, making the response effectively nonlinear. In the steady-state regime, the final output temperature is given by β_z^∞ given by a nonlinear function of β_v (see Eqs. 35 and 36). The machine can implement any linearly separable binary function by appropriately setting the parameters: the qubit energies, the interaction Hamiltonian, and the temperatures of the reference heat baths. Notably, this machine is closely connected to the perceptron model shown in (B), which is used extensively in machine learning. Given inputs x_k , the perceptron first computes a weighted sum y then processed via a nonlinear activation (sigmoid) function f . Similarly, the thermodynamic neuron first creates a virtual qubit at temperature β_v , which is a weighted sum of the input temperatures β_k . Second, the modulator implements the nonlinear activation function. Note that, in a specific regime (ϵ_z sufficiently small), the thermodynamic neuron implements a perceptron as the activation function tends to a sigmoid in this case.

The design of the machine is then completely characterized by the vector \mathbf{h} and the energy gaps ϵ_i for $i \in \{0, 1, \dots, n\}$. These parameters can be chosen freely—they specify the binary function implemented by the thermodynamic neuron.

Let us now discuss the dynamics of the thermodynamic neuron. The machine engineers a virtual qubit at the desired temperature and places it in thermal contact with the output qubit \mathcal{C}_z in resonance with the virtual qubit ($\epsilon_z = \epsilon_v$). This thermal contact is realized via an interaction Hamiltonian $H_{\text{int}} := g(|0\rangle\langle 1|_v \otimes |1\rangle\langle 0|_{\mathcal{C}_z} + \text{h.c.})$. In turn, qubit \mathcal{C}_z thermalizes the output reservoir \mathcal{B}_z to the virtual temperature.

To characterize the virtual temperature, observe that the excited-state population of the virtual qubit in the steady-state reads

$$g_v(\beta_v) := g_0(h_0) \cdot g_1(h_1) \cdot \dots \cdot g_n(h_n) \quad (33)$$

where $g_i(0) = (1 + e^{-\beta_i \epsilon_i})^{-1}$ and $g_i(1) = 1 - g_i(0)$. The virtual temperature β_v satisfies $\exp[-\beta_v \epsilon_v] = g_v(\beta_v)/[1 - g_v(\beta_v)]$ and is given by (see the Supplementary Materials B)

$$\beta_v = \frac{1}{\epsilon_z} \sum_{i=0}^n (-1)^{h_i} \beta_i \epsilon_i \quad (34)$$

The virtual temperature is a linear combination of input temperatures β_i , with relative weights specified by energy gaps ϵ_i and h_i . This relation will be crucial in the next subsection where we establish a connection with perceptrons.

Thermodynamic neuron, in analogy with the inverting thermal machine, features two natural timescales: Thermalization within the collector and the modulator happens quickly, while the thermalization of the output environment \mathcal{B}_z happens slowly. In particular, the time evolution of the output $\beta_z(t)$ is governed by the slow dynamics and given by Eq. 20.

To solve for the inverse steady-state temperature β_z^∞ , we proceed as before (see Eq. 21). We find

$$\beta_z^\infty = \frac{1}{\epsilon_z} \log \left[Q(\beta_v)^{-1} - 1 \right] \quad (35)$$

with $Q(\beta_v) := g_z(\beta_{\text{hot}})g_z(\beta_v) + g_z(\beta_{\text{cold}})[1 - g_z(\beta_v)]$ and β_v is given by Eq. 34. Recall that temperatures β_{cold} and β_{hot} specify the desired temperature range for the computation. Following the derivation in the previous section, we can expand β_z^∞ in the energy gap ϵ_z and obtain

$$\beta_z^\infty = f(\beta_v) + O(\epsilon_z) \quad (36)$$

where $f(x) = (1 + e^x)^{-1}$. Therefore, we see that, for small ϵ_z , the output temperature β_z^∞ behaves essentially as the sigmoid function. For larger values of ϵ_z , the function differs from the sigmoid one but still offers a similar qualitative behavior.

It is important to emphasize that the more inputs a thermodynamic neuron has, the lower is the probability of occupying its virtual subspace. This means that the time it takes to equilibrate the target qubit \mathcal{C}_z to the virtual temperature β_v increases with the number of inputs. To address this challenge, using multiple, interconnected thermodynamic neurons arranged in a network might be more efficient than using a single, complex neuron with many inputs. We will explore how to build such networks of thermodynamic neurons in Network of thermodynamic neurons.

Connection with perceptrons

At this point, it is insightful to establish a formal connection between our model of the thermodynamic neuron and the perceptron

(59). The latter represents the most common model of an artificial neuron and serves as a fundamental component of artificial neural networks.

The perceptron (see Fig. 5B) is a simple algorithm for linear binary classification (60). For a vector of inputs $\mathbf{x} = (x_0, \dots, x_n)$, it produces an output z given by

$$z = f(y) \quad \text{with} \quad y = \sum_{i=0}^n x_i w_i \quad (37)$$

where $x_0 = 1$ by convention, $\mathbf{w} = (w_0, \dots, w_n)$ is a vector of weights that specifies the behavior of the perceptron, and f is the activation function (sigmoid). The perceptron allows for a classification of the input space into two classes; it provides a linear separation of the inputs depending on the value of the function (0 or 1).

At this point, the connection appears clearly. The thermodynamic neuron computes via a two-step procedure, which is very similar to the perceptron. First, given the inputs (encoded here in the temperatures β_1, \dots, β_n), the collector produces a virtual qubit, whose virtual temperature is given by a weighted sum of the input temperatures, with weights given by the energies ϵ_k see (Eq. 34). This corresponds exactly to the computation of the weighted sum y in the perceptron. Second, through the effect of the modulator, the output response becomes nonlinear, and the final temperature β_z^∞ is given by a nonlinear function of the virtual temperature (see Eq. 35). In particular, in the regime of small ϵ_z , this nonlinear functions becomes the sigmoid, hence corresponding exactly to the case of the perceptron (see Eq. 36). This analogy is important and is further illustrated in Fig. 5.

An interesting insight from this analogy is that it sheds light on the importance of the modulator in our model. If the machine would involve only the collector, then the final output temperature would be simply the virtual temperature, corresponding to a trivial activation function $f(y) = y$ in the perceptron algorithm, which is known to perform poorly in machine learning. The modulator provides the essential ingredient of nonlinearity: Its effect is to map the virtual temperature in a nonlinear manner to a temperature inside the range from β_{hot} to β_{cold} . Depending on the value of ϵ_z and the choice of the thermalization model, we get different types of nonlinear function. In particular, when ϵ_z is small and thermalization is a reset model, we get the sigmoid function as in a perceptron. This suggests that thermodynamic neurons could serve as a physical model for a fully analog implementation of perceptrons.

Algorithm for designing the machine

Beyond the conceptual interest, the above connection between the perceptron and our thermodynamic neuron is useful. Suppose we want to design a thermodynamic neuron implementing a given logic operation (e.g., the majority). For this, one would need to find an appropriate combination of qubit energies $\{\epsilon_i\}$ and the vector \mathbf{h} that specifies the interaction Hamiltonian H_{int} . This problem is generally hard and would require rather intensive optimization, especially for more complex functions. Finding the appropriate set of parameters is equivalent to answering the following question: How to choose the systems' local and interacting Hamiltonian so that we achieve the desired steady state? In what follows, we will present a neural network-inspired algorithm, which answers this question quickly and efficiently by finding both the appropriate energy structure and the interaction Hamiltonian of the thermodynamic neuron. Notably, this structure needs to be set only once and, from now on, the

thermodynamic neuron will serve its purpose (i.e., implement the desired function) without any further need of changing its parameters. The algorithm thus provides a general method for designing a thermodynamic neuron implementing arbitrary linearly separable functions.

The main idea of the algorithm is to first run a classical machine learning algorithm that finds the separating hyperplane for the (linearly separable) binary function that one would like to implement. Then, exploiting the formal connection between the perceptron and thermodynamic neuron, one chooses the parameters of the model so that the virtual temperature directly corresponds to the separating hyperplane found by the machine learning algorithm.

Specifically, suppose we want to implement an n -input binary function $R(\mathbf{x})$, where $\mathbf{x} = (x_1, \dots, x_n)$. First, we define the mapping between logical inputs and outputs and temperatures. The logical inputs and output are denoted with $x_1, \dots, x_n, y \in \{0, 1\}$ and encoded in the inverse temperatures of the respective environments through the following procedure

$$x_i = \begin{cases} 0, \beta_i = \beta_{\text{hot}} \\ 1, \beta_i = \beta_{\text{cold}} \end{cases} \quad y = \begin{cases} 0, \beta_z^\infty \leq (1 + \delta)\beta_{\text{hot}} \\ 1, \beta_z^\infty \geq (1 - \delta)\beta_{\text{cold}} \\ \emptyset \quad \text{otherwise} \end{cases} \quad (38)$$

where $i \in \{1, \dots, n\}$. As before, we focus on the range of temperatures from β_{hot} to β_{cold} .

Next, we construct a thermodynamic neuron implementing $R(\mathbf{x})$. For this, we must appropriately set the parameters of the machine, namely, β_0, ϵ_k for $k \in \{0, 1, \dots, n\}$, and the vector \mathbf{h} . Moreover, we also introduce a parameter $\alpha > 0$, which quantifies the overall energy scale of the qubits comprising the machine and hence as well the quality of implementing the desired function. For that, we can use the following algorithm.

Algorithm 1: Designing the thermodynamic neuron

Input: $n, R(\mathbf{x}), \epsilon_z, \alpha$

Output: β_0, ϵ_k , and h_k for $k \in \{0, 1, \dots, n\}$ [see (32)]

Proceed according to the following steps:

1. Construct a training set $D := \{(\mathbf{x}^{(i)}, y_i)\}_{i=1}^{2^n}$, where $\mathbf{x}^{(i)} = [x_1^{(i)}, \dots, x_n^{(i)}]$ and $y_i = R[\mathbf{x}^{(i)}]$.
2. Train a linear classifier (e.g., a sigmoid perceptron) to classify \mathbf{x}_i into two classes: $y_i = 0$ and $y_i = 1$. This gives a vector of weights $\mathbf{w} = (w_0, \dots, w_n)$.
3. Set the elements of the vector as $\mathbf{h} = (h_0, \dots, h_n)$

$$h_k = \begin{cases} 0 & \text{if } w_k \geq 0 \\ 1 & \text{if } w_k < 0 \end{cases} \quad (39)$$

4. Set qubit energies as ϵ_k as

$$\epsilon_k = \begin{cases} \alpha \left| \epsilon_z - \sum_{k=1}^n w_k \right| & \text{if } k=0 \\ \alpha |w_k| & \text{otherwise} \end{cases} \quad (40)$$

5. Set the bias inverse temperature β_0 as

$$\beta_0 = \frac{|w_0|}{\epsilon_z - \sum_{k=1}^n w_k} \quad (41)$$

To see why the above algorithm works, let us observe that the virtual temperature from Eq. 34 becomes

$$\beta_v = \frac{1}{\epsilon_z} \left[(-1)^{i_0} \beta_0 \epsilon_0 + \sum_{k=1}^n (-1)^{i_k} \beta_k \epsilon_k \right] \quad (42)$$

$$= \frac{\alpha}{\epsilon_z} \left[w_0 + \sum_{k=1}^n w_k \beta_k \right] \quad (43)$$

Using the expansion from Eq. 36, we have

$$\beta_z^\infty = f(x) + \mathcal{O}(\epsilon_z), \quad x = \alpha \left(w_0 + \sum_{i=1}^n w_k \beta_k \right) \quad (44)$$

which is exactly the output of the perceptron algorithm for a sigmoid activation function. This demonstrates that the thermodynamic neuron model can implement all functions that can be realized using a (sigmoid) perceptron, namely, all linearly separable functions. The class of functions that can be implemented with a thermodynamic neuron is strictly larger than the sigmoid perceptron, which can be seen by choosing different thermalization models.

Equation 44 also reveals the role of parameter α , which quantifies the steepness of the threshold separating the two outputs or, in other words, the quality of implementing the desired function. In general, α acts as a rescaling of all the energies ϵ_k of the qubits in the collector. Hence, increasing α leads to more dissipation and also lowers the errors in the computation. In particular, for the NOT gate, one can see that $\alpha = \epsilon_1$.

Last, we note that Algorithm 1 should be thought of as a meta-algorithm because it relies on a separate routine to train a linear classifier (step 2). Consequently, its effectiveness and convergence depend on the chosen classifier's properties. Notably, using a classifier with guaranteed convergence translates to similar guarantees for Algorithm 1.

To illustrate how to use Algorithm 1 to design thermodynamic neurons, we now provide two examples.

Example 1: NOR gate

The NOR gate takes $n = 2$ input bits and returns as output the negative OR (see truth table in Fig. 6A). To design the thermodynamic neuron, we follow the steps discussed in Algorithm 1. Using the truth table of NOR, we first construct the set D of $2^n = 4$ data points (see Fig. 6B). In principle, we could now run the algorithm and determine the vector of weights \mathbf{w} . Because, in this case, the separating hyperplane can be found by hand, we simply choose $x_1 + x_2 = 1/2$. This leads to the vector of weights $\mathbf{w} = (1, -2, -2)$. Consequently, the interaction vector \mathbf{h} and energy vector $\epsilon := (\epsilon_0, \epsilon_1, \dots, \epsilon_n)$ become

$$\mathbf{h} = (0, 1, 1), \quad \epsilon = \alpha(\epsilon_z + 4, 2, 2) \quad (45)$$

with the reference (inverse) temperature $\beta_0 = (\epsilon_z + 4)^{-1}$. This choice of parameters leads to the virtual temperature

$$\beta_v = \alpha(1 - 2\beta_1 - 2\beta_2) \quad (46)$$

The machine's response β_z^∞ is then given by Eq. 34 with β_v as given above. In Fig. 6C, we plot the response of the thermodynamic neuron as a function of the input temperatures β_1 and β_2 . The pattern of output temperatures clearly matches the desired NOR function.




Fig. 6. Example 1: NOR. Analysis of the thermodynamic neuron for implementing the NOR function. The truth table of NOR is given in (A). (B) All possible logical states of the machine (blue and red dots) where the color corresponds to the desired output. (C) The response β_z^∞ of the thermodynamic neuron as a function of the inputs β_1 and β_2 . The device does implement the desired NOR gate.

Notably, the NOR function is functionally complete, i.e., any logic function on any number of inputs can be constructed using only NOR functions as building blocks. Consequently, by connecting multiple thermodynamic neurons appropriately, one can, in principle, carry out any classical computation. This shows that the thermodynamic neuron is a universal model of computation.

Example 2: 3-MAJORITY

The 3-majority function takes $n = 3$ inputs bits and outputs most of the input bits. Its truth table is shown in Fig. 7A. To implement 3-MAJORITY using a thermodynamic neuron, we again use Algorithm 1. We construct the training set D of $2^n = 8$ data points (see Fig. 7B). Using the algorithm, we found a vector of weights $\mathbf{w} = (-4, 3, 3, 3)$. The interaction vector \mathbf{h} and the energy vector ϵ are then given by

$$\mathbf{h} = (1, 0, 0, 0), \quad \epsilon = \alpha(\epsilon_z + 12, 3, 3, 3) \quad (47)$$

and the reference temperature is given by $\beta_0 = (\epsilon_z + 12)^{-1}$. This choice of parameters leads to the virtual temperature

$$\beta_v = \alpha(4 - 3\beta_1 - 3\beta_2 - 3\beta_3) \quad (48)$$

As before, the machine's response β_z^∞ is given by Eq. 34 with β_v specified above. In Fig. 7C, we plot the response of the thermodynamic neuron as a function of the input temperatures β_1 , β_2 , and β_3 . The pattern of the output temperatures matches the desired 3-MAJORITY function.

Limitations

From the close connection with perceptrons, we can immediately deduce a general limitation on the class of functions that can be implemented via a single thermodynamic neuron, namely, linearly separable functions.

It is known that perceptron can only represent functions that are linearly separable (61). These are functions for which the set of inputs for which the function takes value 0 can be separated from those whose output is 1 via a simple hyperplane. Consequently, this constraint also limits the range of functions that can be modeled using a single thermodynamic neuron. It is however possible to overcome this limitation by considering networks of neurons. In the next section, we will see how networks of thermodynamic neurons can be used to compute any binary function.

Network of thermodynamic neurons

Perceptrons can be assembled into a network. By increasing the complexity of such a network, it gains the ability to represent more




Fig. 7. Example 2: 3-MAJORITY. Analysis of the thermodynamic neuron for implementing the majority function on three input bits. (A) Truth table. (B) Possible logical states of the machine. The separating hyperplane (dashed line) is specified by the equation $x_1 + x_2 + x_3 = 4/3$. (C) Machine's response β_z^∞ as a function of the inputs β_1 , β_2 , and β_3 . We see that the machine implements the desired operation.

complex functions. According to the universal approximation theorem, a network with sufficiently many layers of perceptrons can approximate any binary function (62). An interesting question is if thermodynamic neurons can also be assembled into networks in such a meaningful manner. In this section, we explore this question in detail.

Combining thermodynamic neurons

In the thermodynamic neuron, the input heat baths are considered to be infinite, while the output heat baths are assumed to have a finite heat capacity. When we connect thermodynamic neurons in a network, the output of some neurons becomes the input for others. However, this poses a challenge: How can we ensure proper functioning of the network when we treat the finite output heat bath of one thermodynamic neuron as the input to another? The finite capacity of the heat bath could disrupt the intended operation of the entire network by introducing unwanted heat currents (e.g., flowing backward). As a result, we can no longer guarantee the validity of Eq. 35 for thermodynamic neurons that constitute the network.

A potential approach to combine thermodynamic neurons is to consider an external agent with access to infinite heat baths at temperatures β_{cold} and β_{hot} . Let us consider a simple network composed of two concatenated thermodynamic neurons. The agent measures the temperature of the output heat bath of the first thermodynamic neuron and, depending on the outcome, couples the input qubit of the second thermodynamic neuron (\mathcal{E}_1) to either β_{cold} or β_{hot} . As a consequence, no unwanted heat currents flow through the output heat bath of the first thermodynamic neuron and the input qubit of the second thermodynamic neuron is coupled to an infinite heat bath.

The proposed method for combining thermodynamic neurons relies on temperature measurements, therefore taking away their autonomy. In the Supplementary Materials C, we present an alternative method of combining thermodynamic neurons that uses a clock. Such a device can be realized autonomously by using an autonomous clock powered by heat baths at different temperatures (42), thus providing a way to make the full computation autonomous (i.e., without invoking external control).

On the basis of the analysis presented above, it is evident that thermodynamic neurons can be interconnected in a manner similar to how perceptrons are linked in artificial neural networks. In this sense, networks composed of thermodynamic neurons can be viewed as analog implementations of neural networks, inheriting the same capacity to perform binary functions. In other words, any

function achievable by a feed-forward neural network can also be realized through a corresponding network of thermodynamic neurons. Given that neural networks are recognized for their ability to approximate any binary function, this implies that networks of thermodynamic neurons can serve as a universal model of computation.

An intriguing direction for further exploration involves considering alternative techniques for connecting thermodynamic neurons that do not necessitate extra thermodynamic resources. Moreover, one could further imagine networks of thermodynamic neurons, which leverage the backflow currents in a useful manner. This could potentially enable feedback within the network, leading to more complex and interesting network dynamics.

Designing networks of thermodynamic neurons

Finding the correct design of a network of thermodynamic neurons for implementing a given function is a nontrivial problem. There are many networks that can implement a given function. Here, we discuss a heuristic approach for determining the network structure for a given binary function. We note that this is only a heuristics, and hence the network of thermodynamic neurons obtained via this method is not guaranteed to implement the correct function.

To find an appropriate set of weights for a network of thermodynamic neurons, we again take inspiration from artificial neural networks. More specifically, suppose we want to implement an n -input binary function $R(\mathbf{x})$. To construct the network implementing $R(\mathbf{x})$, we first choose the structure of the network, i.e., the number of layers, the number of thermodynamic neurons in each layer, and specify the connectivity between thermodynamic neurons. Next, we appropriately choose the free parameters of each of thermodynamic neuron, namely, their reference temperature β_0 , the set of energy gaps $\{\epsilon_k\}$, and the interaction Hamiltonian H_{int} . These parameters can be determined using a straightforward extension of Algorithm 1: Basically, the only difference now is that now the training step (step 2) is performed on the whole network rather than a single thermodynamic neuron. To illustrate this procedure, below we present a network with three thermodynamic neurons for implementing the XOR function, i.e., a function that is not linearly separable.

Example 3: XOR gate

The binary XOR function takes $n = 2$ input bits and returns the parity. It is not a linearly separable function (see Fig. 8B). Hence, it cannot be implemented with a single thermodynamic neuron. To implement XOR, we choose the network structure presented in Fig. 8A. The reason for selecting this particular structure is based on the fact that a binary XOR function can be expressed as a combination of three gates, namely, OR and NAND, whose outputs are fed through an AND gate. The structure of the network we chose mimics this equivalence. Within this network structure, we then use Algorithm 1 to compute the parameters of thermodynamic neurons implementing these three binary functions. Specifically, we construct the corresponding training set D of $2^n = 4$ data points (see Fig. 8B). Then, we perform step 2 of the algorithm using the standard backpropagation algorithm (63) combined with the ADAM optimization (64), obtaining the vectors of weights that correspond to our approximation of the XOR function. Consequently, we use steps 3 to 5 of Algorithm 1 to compute the energy and the interaction vectors as well as the reference bath temperature for each neuron. Thermodynamic neurons are then connected using the method discussed in Combining thermodynamic neurons. The response of the machine, i.e., the inverse temperature of the last thermodynamic




Fig. 8. Example 3: XOR. (A) Structure of a network of thermodynamic neurons that can implement the XOR function. In this case, the training set (i.e., the truth table of the function for all possible inputs) cannot be separated by a hyperplane [see (B)] as the function is not linearly separable. The machine produces the desired response as shown in (C): the response $\beta_z^{(2,1)}$ as a function of the inputs β_1 and β_2 . Note that this machine for implementing XOR can be seen as the composition of a NAND gate and an OR gate, whose outputs $\beta_z^{(1,1)}$ and $\beta_z^{(1,2)}$ are then supplied as an input to an AND gate with output $\beta_z^{(2,1)}$.

neuron, is shown in Fig. 8C as a function of the input temperatures β_1 and β_2 . We see that the network implements the desired XOR function.

DISCUSSION

In this work, we introduced autonomous quantum thermal machines called thermodynamic neurons for performing classical computation. The machine is composed of several qubits, which are coupled to thermal environments at different temperatures. The logical inputs and outputs of the computation are encoded in the temperatures of these environments. By engineering the energies and interactions of the machine's qubits, the device can implement any linearly separable function. In particular, we discussed the implementation of NOT, 3-MAJORITY, and NOR gates, the latter enabling universal computation. For more complex functions, we give an efficient algorithm for tuning the machine parameters. In turn, this algorithm can also be used for networks of thermodynamic neurons, which enable the direct implementation of any desired logical function.

A notable aspect of our machines is that they rely solely on changes in temperature and energy flows: They compute with heat. This sets them apart from conventional (nanoscale) electronic computing devices and other alternative computation models, such as phonon-based computation (65–69), spintronics (70–72), or superconducting circuits (73), where heat-related effects typically hinder computation.

Our work also brings progress from the perspective of autonomous quantum thermal machines by demonstrating a new application for them, namely, classical computation. A single thermodynamic neuron can be considered an autonomous device [see (46)], while their networks can be made autonomous via the addition of a thermodynamic clock (42). An interesting question is if the clock could be directly imbedded in the network of thermodynamic neurons. In parallel, our work also further demonstrates the relevance of virtual qubits and virtual temperatures for computation (29). This complements recent works where these notions are used for characterizing thermodynamic properties of quantum systems (74, 75), the performance of thermal machines (76, 77), and fundamental limits on thermodynamic processes (78).

Another relevant aspect is that our model is thermodynamically consistent, in the sense of complying with the laws of thermodynamics.

This allowed us to investigate its thermodynamic behavior and contrast it with the machine's performance as a computing device. Specifically, for the NOT gate, we observe a clear trade-off between dissipation and performance, in terms of noise robustness. That is, enhancing the performance of the gate requires increasing dissipation. More generally, a similar trade-off relation between dissipation and performance exists for a general computation process carried out by the thermodynamic neuron. It would be interesting to pursue this direction further, e.g., prove a universal relationship by taking inspiration from TURs (79). We want to emphasize that many models of computation consider their thermodynamic aspects under various approximations. Such approximations are generally valid in a specific range of parameters, and outside this range, they can predict an unphysical behavior, e.g., leading to violations of thermodynamic laws. With the growing interest in energy-efficient computing, developing thermodynamically consistent models of computation nowadays becomes increasingly important and has the potential for practical applications.

Outlook

Our work also opens interesting questions from the point of view of machine learning and, more generally, for thermodynamic computing.

As we discussed, thermodynamic neurons have a direct connection to perceptrons and neural networks. In particular, a physical implementation of thermodynamic neurons (and more generally their networks) would provide an alternative physics-based approach for realizing neural networks. This would represent a direct (analog) implementation, hence possibly bypassing some of the challenges of more standard digital (transistor-based) simulations of neural networks. Notably, the energy requirements and heat dissipation of the latter is very substantial, and looking for analog implementations for reducing this thermodynamic cost is important [see, e.g., (80)]. While the current model of thermodynamic neuron is abstract and its potential thermodynamic benefits in comparison to traditional neural network implementations are not yet well understood, investigating the relevance of the thermodynamic neurons in this context is an interesting question.

From a more fundamental perspective, our model could also be used to investigate the thermodynamics of autonomous learning, e.g., using the techniques of refs. (81–83) to modify qubit energies based on the outcome of the computation. In this way, the machine would be able to “learn” a desired behavior in a fully autonomous manner, i.e., to improve its own decisions based on reward or penalty. We believe that this provides an interesting approach for modeling the process of learning in a thermodynamically consistent manner.

Our work can also be discussed from the perspective of thermodynamic computation (19, 20, 27). Here, we believe that an interesting aspect of our model is the fact that computations are implemented in a physical process that is far out of equilibrium. We use machines connected to multiple environments at different temperatures and consider nonequilibrium steady states. What computational power can we obtain from such a model? While we have seen that it can perform universal classical computation and is also naturally connected to neural networks, a key question is to determine its efficiency (notably in terms of time) for solving relevant classes of problems. For example, could this model provide a speedup compared to classical computers for a relevant class of problems?

The performance of thermodynamic neuron depends on how quickly it reaches equilibrium (thermalization). Our simulations with a single neuron show that complete thermalization is not essential. Notably, the qualitative behavior of the model is similar even if it is allowed to thermalize only partially (so-called transient regime). This opens exciting possibilities for exploiting the transient regime to speed up the operation of thermodynamic neurons. At the same time, full thermalization might become more important when combining multiple neurons together. On top of that, thermalization times generally increase with the number of inputs to the thermodynamic neuron. Thus, in some cases, using a longer network of simpler thermodynamic neurons might be a better choice than a shorter network with more complex ones. This is an interesting trade-off that we leave for further exploration in future research.

These are rather long-term perspectives, and a more pressing one is the potential implementation of thermodynamic neurons. In this respect, recent progress on realizing autonomous quantum thermal machines with trapped ions (47) and superconducting qubits (48), together with theoretical proposals in quantum dots (35) and cavity quantum electrodynamics (37) are relevant. An interesting alternative is to investigate if the physics of our model can be reproduced by a fully classical model based on rate equations. This would open the door to a classical implementation within stochastic thermodynamics (84).

Supplementary Materials

This PDF file includes:

Supplementary Text

Fig. S1

REFERENCES AND NOTES

1. C. H. Bennett, The thermodynamics of computation—A review. *Int. J. Theor. Phys.* **21**, 905–940 (1982).
2. J. M. R. Parrondo, J. M. Horowitz, T. Sagawa, Thermodynamics of information. *Nat. Phys.* **11**, 131–139 (2015).
3. D. H. Wolpert, The stochastic thermodynamics of computation. *J. Phys. A Math. Theor.* **52**, 193001–193114 (2019).
4. S. Deffner, C. Jarzynski, Information processing and the second law of thermodynamics: An inclusive, Hamiltonian approach. *Phys. Rev. X* **3**, 041003 (2013).
5. A. B. Boyd, D. Mandal, J. P. Crutchfield, Thermodynamics of modularity: Structural costs beyond the Landauer bound. *Phys. Rev. X* **8**, 031036–031058 (2018).
6. P. Faist, R. Renner, Fundamental work cost of quantum processes. *Phys. Rev. X* **8**, 021011–021030 (2018).
7. J. Gu, P. Gaspard, Microreversibility, fluctuations, and nonlinear transport in transistors. *Phys. Rev. E* **99**, 012137–012153 (2019).
8. D. H. Wolpert, A. Kolchinsky, Thermodynamics of computing with circuits. *New J. Phys.* **22**, 063047–063053 (2020).
9. C. Y. Gao, D. T. Limmer, Principles of low dissipation computing from a stochastic circuit model. *Phys. Rev. Res.* **3**, 033169–033187 (2021).
10. N. Feistas, J.-C. Delvenne, M. Esposito, Stochastic thermodynamics of nonlinear electronic circuits: A realistic framework for computing around kT . *Phys. Rev. X* **11**, 031064–031091 (2021).
11. P. Helms, D. T. Limmer, Stochastic thermodynamic bounds on logical circuit operation. arXiv: 2211.00670 [cond-mat.stat-mech] (2022).
12. J. Kuang, X. Ge, Y. Yang, L. Tian, Modeling and optimization of low-power and gates based on stochastic thermodynamics. *IEEE Trans. Circuits Syst. II Express Briefs* **69**, 3729–3733 (2022).
13. U. Seifert, Stochastic thermodynamics, fluctuation theorems and molecular machines. *Rep. Prog. Phys.* **75**, 126001–1260106 (2012).
14. A. Solfanelli, A. Santini, M. Campisi, Quantum thermodynamic methods to purify a qubit on a quantum processing unit. *AVS Quantum Sci.* **4**, 026802–026809 (2022).
15. M. Fellous-Asiani, J. H. Chai, R. S. Whitney, A. Auffèves, H. K. Ng, Limitations in quantum computing from resource constraints. *PRX Quantum* **2**, 040335–040346 (2021).

16. M. Fellous-Asiani, J. H. Chai, Y. Thonnart, H. K. Ng, R. S. Whitney, A. Auffèves, Optimizing resource efficiencies for scalable full-stack quantum computers. *arXiv*: 2209.05469 [quant-ph] (2022).
17. A. Auffèves, Quantum technologies need a quantum energy initiative. *PRX Quantum* **3**, 020101–020113 (2022).
18. J. Stevens, D. Szombati, M. Maffei, C. Elouard, R. Assouly, N. Cottet, R. Dassonneville, Q. Ficheux, S. Zeppetzauer, A. Bienfait, A. N. Jordan, A. Auffèves, B. Huard, Energetics of a single qubit gate. *Phys. Rev. Lett.* **129**, 110601 (2022).
19. T. Conte, E. De Benedictis, N. Ganesh, T. Hylton, J. P. Strachan, R. S. Williams, A. Alemi, L. Altenberg, G. Crooks, J. Crutchfield, L. del Rio, J. Deutsch, M. De Weese, K. Douglas, M. Esposito, M. Frank, R. Fry, P. Harsha, M. Hill, C. Kello, J. Krichmar, S. Kumar, S.-C. Liu, S. Lloyd, M. Marsili, I. Nemenman, A. Nugent, N. Packard, D. Randall, P. Sadowski, N. Santhanam, R. Shaw, A. Stieg, E. Stopnitzky, C. Teuscher, C. Watkins, D. Wolpert, J. Yang, Y. Yufik, Thermodynamic computing. *arXiv*: 1911.01968 [cs.CY] (2019).
20. P. J. Coles, C. Szczepanski, D. Melanson, K. Donatella, A. J. Martinez, F. Sbahi, Thermodynamic AI and the fluctuation frontier. *arXiv*: 2302.06584 [cs.ET] (2023).
21. M. Aifer, D. Melanson, K. Donatella, G. Crooks, T. Ahle, P. J. Coles, Error mitigation for thermodynamic computing. *arXiv*: 2401.16231 [cs.ET] (2024).
22. S. Duffield, M. Aifer, G. Crooks, T. Ahle, P. J. Coles, Thermodynamic matrix exponentials and thermodynamic parallelism. *arXiv*: 2311.12759 [cond-mat.stat-mech] (2023).
23. S. Goldt, U. Seifert, Stochastic thermodynamics of learning. *Phys. Rev. Lett.* **118**, 010601 (2017).
24. T. Hylton, Thermodynamic neural network. *Entropy* **22**, 256–280 (2020).
25. T. Hylton, Thermodynamic state machine network. *Entropy* **24**, 744–769 (2022).
26. A. B. Boyd, J. P. Crutchfield, M. Gu, Thermodynamic machine learning through maximum work production. *New J. Phys.* **24**, 083040–083076 (2022).
27. M. Aifer, K. Donatella, M. H. Gordon, S. Duffield, T. Ahle, D. Simpson, G. E. Crooks, P. J. Coles, Thermodynamic linear algebra. *arXiv*: 2308.05660 [cond-mat.stat-mech] (2023).
28. N. Linden, S. Popescu, P. Skrzypczyk, How small can thermal machines be? The smallest possible refrigerator. *Phys. Rev. Lett.* **105**, 130401–130407 (2010).
29. N. Brunner, N. Linden, S. Popescu, P. Skrzypczyk, Virtual qubits, virtual temperatures, and the foundations of thermodynamics. *Phys. Rev. E* **85**, 051117–051133 (2012).
30. J. Goold, M. Huber, A. Riera, L. Rio, P. Skrzypczyk, The role of quantum information in thermodynamics—A topical review. *J. Phys. A* **49**, 143001–143035 (2016).
31. S. Vinjanampathy, J. Anders, Quantum thermodynamics. *Contemp. Phys.* **57**, 545–579 (2016).
32. M. T. Mitchison, Quantum thermal absorption machines: Refrigerators, engines and clocks. *Contemp. Phys.* **60**, 164–187 (2019).
33. H. E. D. Scovil, E. O. Schulz-DuBois, Three-level masers as heat engines. *Phys. Rev. Lett.* **2**, 262–263 (1959).
34. A. Levy, R. Kosloff, Quantum absorption refrigerator. *Phys. Rev. Lett.* **108**, 070604–070609 (2012).
35. D. Venturelli, R. Fazio, V. Giovannetti, Minimal self-contained quantum refrigeration machine based on four quantum dots. *Phys. Rev. Lett.* **110**, 256801–256806 (2013).
36. L. A. Correa, J. P. Palao, D. Alonso, G. Adesso, Quantum-enhanced absorption refrigerators. *Sci. Rep.* **4**, 3949–3958 (2014).
37. P. P. Hofer, J. R. Souquet, A. A. Clerk, Quantum heat engine based on photon-assisted cooper pair tunneling. *Phys. Rev. B* **93**, 041418–041423 (2016).
38. D. Gelbwaser-Klimovsky, R. Alicki, G. Kurizki, Minimal universal quantum heat machine. *Phys. Rev. E* **87**, 012140–012149 (2013).
39. P. Strasberg, C. W. Wächter, G. Schaller, Autonomous implementation of thermodynamic cycles at the nanoscale. *Phys. Rev. Lett.* **126**, 180605–180611 (2021).
40. W. Niedenzu, M. Huber, E. Boukobza, Concepts of work in autonomous quantum heat engines. *Quantum* **3**, 195 (2019).
41. J. B. Brask, N. Brunner, G. Haack, M. Huber, Autonomous quantum thermal machine for generating steady-state entanglement. *New J. Phys.* **17**, 113029–113039 (2015).
42. P. Erker, M. T. Mitchison, R. Silva, M. P. Woods, N. Brunner, M. Huber, Autonomous quantum clocks: Does thermodynamics limit our ability to measure time? *Phys. Rev. X* **7**, 031022–031034 (2017).
43. E. Schwarzhans, M. P. E. Lock, P. Erker, N. Friis, M. Huber, Autonomous temporal probability concentration: Clockworks and the second law of thermodynamics. *Phys. Rev. X* **11**, 011046–011068 (2021).
44. M. P. Woods, Autonomous ticking clocks from axiomatic principles. *Quantum* **5**, 381 (2021).
45. P. P. Hofer, J. B. Brask, M. Perarnau-Llobet, N. Brunner, Quantum thermal machine as a thermometer. *Phys. Rev. Lett.* **119**, 090603–090610 (2017).
46. J. A. Marín Guzmán, P. Erker, S. Gasparinetti, M. Huber, N. Yunger Halpern, Divincenzo-like criteria for autonomous quantum machines. *arXiv*: 2307.08739 [quant-ph] (2023).
47. G. Maslennikov, S. Ding, R. Hablützel, J. Gan, A. Roulet, S. Nimmrichter, J. Dai, V. Scarani, D. Matsukevich, Quantum absorption refrigerator with trapped ions. *Nat. Commun.* **10**, 202–210 (2019).
48. M. A. Amir, P. J. Suria, J. A. Marín Guzmán, C. Castillo-Moreno, J. M. Epstein, N. Yunger Halpern, S. Gasparinetti, Thermally driven quantum refrigerator autonomously resets superconducting qubit. *arXiv*: 2305.16710 [quant-ph] (2023).
49. M. P. Woods, M. Horodecki, Autonomous quantum devices: When are they realizable without additional thermodynamic costs? *Phys. Rev. X* **13**, 011016–011047 (2023).
50. P. P. Hofer, M. Perarnau-Llobet, L. D. M. Miranda, G. Haack, R. Silva, J. B. Brask, N. Brunner, Markovian master equations for quantum thermal machines: Local versus global approach. *New J. Phys.* **19**, 123037–123058 (2017).
51. R. C. Tolman, P. C. Fine, On the irreversible production of entropy. *Rev. Mod. Phys.* **20**, 51–77 (1948).
52. G. T. Landi, M. Paternostro, Irreversible entropy production: From classical to quantum. *Rev. Mod. Phys.* **93**, 035008–035067 (2021).
53. J. M. Horowitz, T. R. Gingrich, Proof of the finite-time thermodynamic uncertainty relation for steady-state currents. *Phys. Rev. E* **96**, 020103–020107 (2017).
54. A. C. Barato, U. Seifert, Thermodynamic uncertainty relation for biomolecular processes. *Phys. Rev. Lett.* **114**, 158101 (2015).
55. G. Falasco, M. Esposito, J.-C. Delvenne, Unifying thermodynamic uncertainty relations. *New J. Phys.* **22**, 053046–053062 (2020).
56. N. Shiraishi, K. Funo, K. Saito, Speed limit for classical stochastic processes. *Phys. Rev. Lett.* **121**, 070601–070607 (2018).
57. B. Sothmann, R. Sánchez, A. N. Jordan, Thermoelectric energy harvesting with quantum dots. *Nanotechnology* **26**, 032001–032047 (2015).
58. D. Janzing, P. Wocjan, R. Zeier, R. Geiss, T. Beth, Thermodynamic cost of reliability and low temperatures: Tightening Landauer's principle and the second law. *Int. J. Theor. Phys.* **39**, 2717–2753 (2000).
59. W. S. McCulloch, W. Pitts, A logical calculus of the ideas immanent in nervous activity. *Bull. Math. Biophys.* **5**, 115–133 (1943).
60. F. Rosenblatt, *The Perceptron, a Perceiving and Recognizing Automaton Project Para* (Cornell Aeronautical Laboratory, 1957).
61. I. Goodfellow, Y. Bengio, A. Courville, *Deep Learning* (MIT Press, 2016).
62. K. Hornik, M. Stinchcombe, H. White, Multilayer feedforward networks are universal approximators. *Neural Netw.* **2**, 359–366 (1989).
63. D. E. Rumelhart, G. E. Hinton, R. J. Williams, Learning representations by back-propagating errors. *Nature* **323**, 533–536 (1986).
64. D. P. Kingma, J. Ba, Adam: A method for stochastic optimization. *arXiv*: 1412.6980 [cs.LG] (2017).
65. R. Ruskov, C. Tahan, On-chip cavity quantum phonodynamics with an acceptor qubit in silicon. *Phys. Rev. B* **88**, 064308–064315 (2013).
66. S. R. Sklan, Splash, pop, sizzle: Information processing with phononic computing. *AIP Adv.* **5**, 053302 (2015).
67. M.-A. Lemonde, S. Meesala, A. Sipahigil, M. J. A. Schuetz, M. D. Lukin, M. Loncar, P. Rabl, Phonon networks with silicon-vacancy centers in diamond waveguides. *Phys. Rev. Lett.* **120**, 213603–213610 (2018).
68. M. V. Gustafsson, T. Aref, A. F. Kockum, M. K. Ekström, G. Johansson, P. Delsing, Propagating phonons coupled to an artificial atom. *Science* **346**, 207–211 (2014).
69. W. Chen, Y. Lu, S. Zhang, K. Zhang, G. Huang, M. Qiao, X. Su, J. Zhang, J. Zhang, L. Banchi, M. S. Kim, K. Kim, Scalable and programmable phononic network with trapped ions. *Nat. Phys.* **2023**, 1–7 (2023).
70. S. A. Wolf, A. Y. Chelkanova, D. M. Treger, Spintronics—A retrospective and perspective. *IBM J. Res. Dev.* **50**, 101–110 (2006).
71. A. Mahmoud, F. Ciubotaru, F. Vandeveken, A. V. Chumak, S. Hamdioui, C. Adelman, S. Cotofana, Introduction to spin wave computing. *J. Appl. Phys.* **128**, 161101–161142 (2020).
72. S. K. Kim, G. S. D. Beach, K. J. Lee, T. Ono, T. Rasing, H. Yang, Ferrimagnetic spintronics. *Nat. Mat.* **21**, 24–34 (2022).
73. C. Z. Pratt, K. J. Ray, J. P. Crutchfield, Dynamical computing on the nanoscale: Superconducting circuits for thermodynamically-efficient classical information processing. *arXiv*: 2307.01926 [cond-mat.stat-mech] (2023).
74. P. Skrzypczyk, R. Silva, N. Brunner, Passivity, complete passivity, and virtual temperatures. *Phys. Rev. E* **91**, 052133–052137 (2015).
75. P. Lipka-Bartosik, P. Skrzypczyk, All states are universal catalysts in quantum thermodynamics. *Phys. Rev. X* **11**, 011061–011091 (2021).
76. R. Silva, G. Manzano, P. Skrzypczyk, N. Brunner, Performance of autonomous quantum thermal machines: Hilbert space dimension as a thermodynamical resource. *Phys. Rev. E* **94**, 032120–032135 (2016).
77. A. Usui, W. Niedenzu, M. Huber, Simplifying the design of multilevel thermal machines using virtual qubits. *Phys. Rev. A* **104**, 042224–042238 (2021).
78. F. Clivaz, R. Silva, G. Haack, J. B. Brask, N. Brunner, M. Huber, Unifying paradigms of quantum refrigeration: A universal and attainable bound on cooling. *Phys. Rev. Lett.* **123**, 170605–170611 (2019).
79. U. Seifert, From stochastic thermodynamics to thermodynamic inference. *Annu. Rev. Condens. Matter Phys.* **10**, 171–192 (2019).
80. H. Wang, Analog chip paves the way for sustainable AI. *Nature* **620**, 731–732 (2023).
81. N. C. Keim, J. D. Paulsen, Z. Zervavic, S. Sastry, S. R. Nagel, Memory formation in matter. *Rev. Mod. Phys.* **91**, 035002–035028 (2019).

82. V. Lopez-Pastor, F. Marquardt, Self-learning machines based on Hamiltonian echo backpropagation. *Phys. Rev. X* **13**, 031020–031054 (2023).
83. W. Zhong, J. M. Gold, S. Marzen, J. L. England, N. Yunger Halpern, Machine learning outperforms thermodynamics in measuring how well a many-body system learns a drive. *Sci. Rep.* **11**, 9333–9344 (2021).
84. S. Ciliberto, Experiments in stochastic thermodynamics: Short history and perspectives. *Phys. Rev. X* **7**, 021051–021077 (2017).

Acknowledgments: We are grateful to G. Haack and P. Skrzypczyk for fruitful discussions. We also thank N. Yunger Halpern, M. Huber, J. A. Marín Guzmán, and P. Coles for useful comments on the first draft of this paper. **Funding:** We acknowledge the Swiss National Science Foundation for financial support through the Ambizione grant PZ00P2-186067 and the NCCR

SwissMAP. **Author contributions:** Conceptualization, methodology, validation, formal analysis, investigation, writing, and visualization: P.L.-B. Conceptualization, methodology, validation, investigation, writing, and supervision: M.P.-L. Conceptualization, methodology, validation, investigation, writing, supervision, and funding acquisition: N.B. **Competing interests:** The authors declare that they have no competing interests. **Data and materials availability:** All data needed to evaluate the conclusions in the paper are present in the paper and/or the Supplementary Materials.

Submitted 10 November 2023

Accepted 30 July 2024

Published 4 September 2024

10.1126/sciadv.adm8792